

REPORT DOCUMENTATION PAGE

Form Approved
OMB NO. 0704-0188

Public Reporting burden for this collection of information is estimated to average 1 hour per response, including the time for reviewing instructions, searching existing data sources, gathering and maintaining the data needed, and completing and reviewing the collection of information. Send comment regarding this burden estimate or any other aspect of this collection of information, including suggestions for reducing this burden, to Washington Headquarters Services, Directorate for Information Operations and Reports, 1215 Jefferson Davis Highway, Suite 1204, Arlington, VA 22202-4302, and to the Office of Management and Budget, Paperwork Reduction Project (0704-0188), Washington, DC 20503.

1. AGENCY USE ONLY (Leave Blank)		2. REPORT DATE March 28, 1999	3. REPORT TYPE AND DATES COVERED Final Progress
4. TITLE AND SUBTITLE SCALABLE PARALLEL ALGORITHMS		5. FUNDING NUMBERS DAAH04-95-1-0538	
6. AUTHOR(S) Vipin Kumar			
7. PERFORMING ORGANIZATION NAME(S) AND ADDRESS(ES) Computer Science & Engineering Dept., University of MN 200 Union St. SE, Minneapolis, MN 55455-0159		8. PERFORMING ORGANIZATION REPORT NUMBER	
9. SPONSORING / MONITORING AGENCY NAME(S) AND ADDRESS(ES) U. S. Army Research Office P.O. Box 12211 Research Triangle Park, NC 27709-2211		10. SPONSORING / MONITORING AGENCY REPORT NUMBER ARO 33984.3-MA	
11. SUPPLEMENTARY NOTES The views, opinions and/or findings contained in this report are those of the author(s) and should not be construed as an official Department of the Army position, policy or decision, unless so designated by the documentation.			
12 a. DISTRIBUTION / AVAILABILITY STATEMENT Approved for public release; distribution unlimited.		12 b. DISTRIBUTION CODE	
13. ABSTRACT (Maximum 200 words) The objective of this project is to develop scalable parallel formulations of the key computational kernels used in scientific simulations. The specific problems investigated in this project are fast and high quality graph partitioners, highly parallel direct solvers, and parallel formulations of robust preconditioners for iterative solvers, as well as parallel formulations of particle simulation techniques. We have developed a fast and high quality parallel formulations of our multilevel graph partitioning algorithm that are able to partition very large graphs quickly on parallel computers, making it feasible to perform frequent repartitioning of the adaptive and unstructured mesh in adaptive FEM computations. We have developed massively parallel formulations of particle simulation techniques such as Fast Multipole and Barnes-Hut methods, and have investigated the use of this formulation for solving dense linear systems arising in boundary element solution of integral equations. We have developed an MPI-based portable library, called PSPASES, that has been used to solve some of the largest sparse linear systems that have been solved using direct methods. We have also developed robust and parallel preconditioners for iterative solvers using our fast graph partitioning technique and highly parallel Cholesky factorization.			
14. SUBJECT TERMS Parallel Algorithms, Sparse Linear System Solver, Graph Partitioning, Cholesky Factorization		15. NUMBER OF PAGES	
		16. PRICE CODE	
17. SECURITY CLASSIFICATION OR REPORT UNCLASSIFIED	18. SECURITY CLASSIFICATION ON THIS PAGE UNCLASSIFIED	19. SECURITY CLASSIFICATION OF ABSTRACT UNCLASSIFIED	20. LIMITATION OF ABSTRACT UL

19990816 028

SCALABLE PARALLEL ALGORITHMS

Final Report

Vipin Kumar

**April 8, 1999
U. S. Army Research Office**

Contract / Grant Number: DA/DAAH04-95-1-0538

**Computer Science Department
University of Minnesota
Minneapolis, MN 55455**

**Approved for Public Release;
Distribution Unlimited.**

Contents

1 Problem Statement	4
2 Summary of Important Results	4
3 List of Publications Resulting from the Grant	6
4 List of Participating Personnel	8

Abstract

The objective of this project is to develop scalable parallel formulations of the key computational kernels used in scientific simulations. The specific problems investigated in this project are fast and high quality graph partitioners, highly parallel direct solvers, and parallel formulations of robust preconditioners for iterative solvers, as well as parallel formulations of particle simulation techniques. We have developed a fast and high quality parallel formulations of our multilevel graph partitioning algorithm that are able to partition very large graphs quickly on parallel computers, making it feasible to perform frequent repartitioning of the adaptive and unstructured mesh in adaptive FEM computations. We have developed massively parallel formulations of particle simulation techniques such as Fast Multipole and Barnes-Hut methods, and have investigated the use of this formulation for solving dense linear systems arising in boundary element solution of integral equations. We have developed an MPI-based portable library, called PSPASES, that has been used to solve some of the largest sparse linear systems that have been solved using direct methods. We have also developed robust and parallel preconditioners for iterative solvers using our fast graph partitioning technique and highly parallel Cholesky factorization.

1 Problem Statement

Virtually all scientific and natural phenomena can be modeled as systems of differential equations that are solved using finite element and finite difference methods. The objective of this project is to solve linear systems arising from these methods. These sparse linear systems are too large to be solved cost effectively on traditional vector-supercomputers. This project aims at developing highly parallel linear system solvers and investigating their applications in problems of interest to US Army. This work has considerable significance since it will enable modeling accuracies and discretizations much finer than currently possible. It will also result in robust and portable software that can be used for a variety of applications.

2 Summary of Important Results

We have developed highly efficient parallel formulations of our multi-level graph partitioning algorithm that achieve high degree of concurrency, while maintaining the high quality partitions produced by the serial multi-level algorithm. Our parallel graph partitioning algorithm can partition very large graphs (e.g., over 8 million vertices) in 128 parts on a 128-processor Cray T3D in under 10 seconds. Thus, our parallel algorithm makes it possible to perform frequent mesh re-partitioning in adaptive computations without compromising quality.

In adaptive finite element computation, dynamic adjustments to the mesh require repartitioning of the mesh to improve load balance. This re-partitioning also results in movement of data structures associated with graph vertices. Hence, a good re-partitioning algorithm should minimize the movement of vertices (in addition to balancing the load and minimizing the cut of the resulting new partition). In this project, we have developed parallel graph partitioners that minimize the movement of data (in addition to minimizing the edge cut of the partitioning) when a mesh is adaptively refined. This partitioner will also facilitate the development of highly parallel mesh generators, as the mesh generation algorithms also generate the mesh by adaptively refining it at various places.

Factorization algorithms based on threshold incomplete LU factorization (ILUT) have been found to be quite effective in preconditioning iterative system solvers. However, because these factorizations allow the fill elements to be created dynamically, their parallel formulations had not been well understood, and they had been considered to be unsuitable for distributed memory parallel computers. We have developed a highly parallel formulation of the ILUT factorization algorithm for distributed memory parallel computers. This algorithm uses our graph partitioning algorithm in conjunction with a parallel maximal independent subset algorithm to effectively parallelize both the factorization as well as the solution of the resulting triangular factors. Our experiments have shown that both the ILUT factorization as well as the solution of the resulting triangular systems can be performed very fast on distributed memory parallel computers. Furthermore, our experiments using the GMRES iterative solver show that the amount of time spent in computing the factorization is usually much less than the amount of time required to solve the systems.

We have developed massively parallel formulations of particle simulation techniques such as

Fast Multipole and Barnes-Hut methods, and have used this formulation for solving dense linear systems arising in boundary element solution of integral equations. For a problem with 200,000 unknowns, we demonstrate over two orders of magnitude speedup from parallelization and another two orders of magnitude from approximation in our preliminary implementation.

3 List of Publications Resulting from the Grant

1. Ananth Grama, Vipin Kumar and Ahmed Sameh, "Parallel Iterative Solvers and Preconditioners Using Approximate Hierarchical Methods", To be presented at the Copper Mountain Conference on Iterative Methods, April 1996, Copper Mountain, CO.
2. George Karypis and Vipin Kumar, Parallel Multi-level Graph Partitioning. Proceedings of the International Parallel Processing Symposium, April 1996.
3. Anshul Gupta, George Karypis and Vipin Kumar, A Highly Scalable Parallel Algorithm for Sparse Matrix Factorization, IEEE Transactions on Parallel and Distributed Systems Volume 8, Number 5, May 1997. A short version of this paper won the Outstanding Student Paper Award from the Supercomputing 94 conference.
4. George Karypis and Vipin Kumar, Parallel Multi-level k-way Partitioning Scheme for Irregular Graphs, Proceedings of Supercomputing'96, Pittsburgh, November 1996.
5. A. Grama, V. Kumar, and A. Sameh, Parallel Hierarchical Solvers and Preconditioners for Boundary Element Methods, Proceedings of Supercomputing'96, Pittsburgh, November 1996. Best Student paper Nominee.
6. Ananth Grama, Vipin Kumar, Sanjay Ranka, and Vineet Singh, A3: A Simple and Asymptotically Accurate Model for Parallel Computation, Proceedings of the Sixth Symposium on Frontiers of Massively Parallel Computing, Annapolis, MD, October 1996.
7. Anshul Gupta, Fred Gustavson, Mahesh Joshi, George Karypis, Vipin Kumar, Design and Implementation of a Scalable Parallel Direct Solver for Sparse Symmetric Positive Definite Systems. Proceedings of the Eight SIAM Conference Conference on Parallel processing, March 1997.
8. V. Kumar, G. Karypis, and A. Grama, Role of Message-Passing in Performance Oriented Parallel Programming, Proceedings of the Eight SIAM Conference Conference on Parallel processing, March 1997.
9. E.H. Han and G. Karypis and V. Kumar, Scalable Parallel Data Mining for Association Rules, Proceedings of 1997 ACM-SIGMOD International Conference on Management of Data, May 1997.
10. G. Karypis and R. Aggarwal and V. Kumar and S. Shekhar, Multilevel Hypergraph Partitioning: Application in VLSI Domain, Proceedings ACM/IEEE Design Automation Conference, June, 1997.
11. Kirk Schloegel, George Karypis and Vipin Kumar, Multilevel Diffusion Schemes for Repartitioning of Adaptive Meshes, Journal of Parallel and Distributed Computing, Volume 47, Number 2, Dec 15, 1997, pp 109-124.

12. George Karypis and Vipin Kumar, A Parallel Algorithm for Multilevel Graph partitioning and Sparse Matrix Ordering, *Journal of Parallel and Distributed Computing*, Volume 48, Number 1, Jan 1998.
13. George Karypis and Vipin Kumar, Multilevel K-way Partitioning Scheme for Irregular Graphs. *Journal of Parallel and Distributed Computing*, Volume 48, Number 1, Jan 1998.
14. Ananth Y. Grama, Vipin Kumar and Ahmed Sameh, Scalable Parallel Formulations of the Barnes-Hut Method for n-Body Simulations, *Parallel Computing*, Volume 24, Number 5-6, June 1998.
15. Ananth Grama, Vipin Kumar and Ahmed Sameh, "Parallel Hierarchical Solvers and Preconditioners for Boundary Element Methods", *SIAM Journal on Scientific Computing*, accepted for publication, 1997.
16. S. Shekhar, S. Ravada, V. Kumar, G. Turner, and D. Chubb, Declustering and Load Balancing methods for Parallelizing Geographical Information Systems, *IEEE Transactions on Knowledge and Data Engineering*, Volume 10, Number 4, 1998.
17. Mahesh V. Joshi, George Karypis, and Vipin Kumar, "ScalParC: A New Scalable and Efficient Parallel Classification Algorithm for Mining Large Datasets", *Proc. of 12th International Parallel Processing Symposium (IPPS98)*, Orlando, March 1998.
18. Eui-Hong (Sam) Han, George Karypis, Vipin Kumar and Bamshad Mobasher, Clustering Based On Association Rule Hypergraphs (1997), *Workshop on Research Issues on Data Mining and Knowledge Discovery*, 1997.
19. Jerome Moore, Eui-Hong (Sam) Han, Daniel Boley, Maria Gini, Robert Gross, Kyle Hastings, George Karypis, Vipin Kumar, and Bamshad Mobasher, Web Page Categorization and Feature Selection Using Association Rule and Principal Component Clustering, *Workshop on Information Technologies and Systems*, 1997.
20. Mahesh V. Joshi, George Karypis, and Vipin Kumar, A High Performance Two Dimensional Scalable Parallel Algorithm for Solving Sparse Triangular Systems, *4th International Conference on High Performance Computing*, December 1997.
21. Kirk Schloegel, George Karypis and Vipin Kumar, Repartitioning of Adaptive Meshes: Experiments with Multilevel Diffusion, *Proceedings of Euro-Par'97*, August 1997.
22. George Karypis and Vipin Kumar, Parallel Threshold-Based ILUT Factorization, *Proceedings of Supercomputing'97*. Extended version available as Tech Report 96-061, department of computer science, University of Minnesota, 1996.

4 List of Participating Personnel

- Vipin Kumar, PI
- Research Assistants: Anurag Srivastava, Kirk Schloegel, Dalvinder Malhotra, Sam Han, Mahesh Joshi
- Research Faculty and Post Doctoral Associates: Ananth Grama, George Karypis

George Karypis and Ananth Grama graduated with a PhD degree in computer science. Tom Nurkkala (supported by a companion Aasert grant) graduated with a PhD in computer science. Anurag Srivastava and Dalvinder Malhotra graduated with M.Sc. degrees in computer science.

The dissertation of Kumar's PhD student Anshul Gupta was selected as one of the five finalists in The ACM Doctoral Dissertation Competition for 1995.

The paper by Kumar's students Anshul Gupta and George Karypis on Scalable Sparse Matrix Factorization won Mannheim SuParCup Prize from the 1995 European Supercomputing conference.

The paper by Kumar's student Ananth Grama (co-authored with Kumar and Sameh), Parallel Hierarchical Solvers and Preconditioners for Boundary Element Methods, was selected as Best Student paper Nominee for Supercomputing'96, Pittsburgh, November 1996.